

# Selecting Bases in Spectral learning of Predictive State Representations via Model Entropy

Yunlong Liu    YLLIU@XMU.EDU.CN and Hexing Zhu    HXZHU\_XMU@FOXMAIL.COM  
*Department of Automation, Xiamen University, Xiamen, China, 361005*

## Abstract

Predictive State Representations (PSRs) are powerful techniques for modelling dynamical systems, which represent a state as a vector of predictions about future observable events (tests). In PSRs, one of the fundamental problems is the learning of the PSR model of the underlying system. Recently, spectral methods have been successfully used to address this issue by treating the learning problem as the task of computing an singular value decomposition (SVD) over a submatrix of a special type of matrix called the Hankel matrix. Under the assumptions that the rows and columns of the submatrix of the Hankel Matrix are sufficient (which usually means a very large number of rows and columns, and almost fails in practice) and the entries of the matrix can be estimated accurately, it has been proven that the spectral approach for learning PSRs is statistically consistent and the learned parameters can converge to the true parameters. However, in practice, due to the limit of the computation ability, only a finite set of rows or columns can be chosen to be used for the spectral learning. While different sets of columns usually lead to variant accuracy of the learned model, in this paper, we propose an approach for selecting the set of columns, namely basis selection, by adopting a concept of model entropy to measure the accuracy of the learned model. Experimental results are shown to demonstrate the effectiveness of the proposed approach.

**Keywords:** Predictive State Representations; Spectral Methods; Basis Selection

## 1. Introduction

Modelling dynamical systems is a common problem in science and engineer, and there are many applications related to such a modelling, e.g. robot navigation, natural language processing, speech recognition, etc. [Hamilton et al. \(2013\)](#). Many approaches have been devoted to modelling systems, such as Hidden Markov Models (HMM) for uncontrolled dynamical systems and Partially Observable Markov Decision Processes (POMDPs) for controlled systems. However, such latent-state approaches usually suffer from local minima and require some assumptions [Singh and James \(2004\)](#). Predictive State Representations (PSRs) offer an effective approach for modelling partially dynamical systems [Littman et al. \(2001\)](#). Unlike the latent-state approaches, PSRs use a vector of predictions about future events, called tests, to represent the state. The tests can be executed on the systems and are fully observable. Compared to the latent-state approaches, PSRs have shown many advantages, such as the possibility of obtaining a global optimal model, more expressive power and less required prior domain knowledge, etc. [Liu et al. \(2015\)](#).

There are two main problems in PSRs. One is the learning of the PSR model; the other is the application of the learned model, including predicting and planning [Rosencrantz](#)

et al. (2004); Liu et al. (2014). The state-of-the-art technique for addressing the learning problem is the spectral approach Boots et al. (2011). Spectral methods treat the learning problem as the task of computing a singular value decomposition (SVD) over a submatrix  $H_s$  of a special type of matrix called the Hankel matrix Hsu et al. (2009). Under the strong assumptions that the rows  $H$  and columns  $T$  of  $H_s$  are sufficient and the entries of the matrix can be estimated accurately, it has been proven that the spectral approach for learning PSRs is statistically consistent and the learned parameters can converge to the true parameters Boots et al. (2011). However, for the sufficient assumption, it usually means a very large number of rows or columns of the  $H_s$  matrix, which almost fails in practice Kulesza et al. (2015). At the same time, the computation complexity of the SVD operation on  $H_s$  is  $O(|T|^2|H|)$  ( $|T|$  and  $|H|$  is the number of the columns and rows in the matrix respectively), for large set of  $T$  and  $H$ , such an operation is also prohibitively expensive. Also, for the spectral approach, to obtain the model parameter, one should estimate and manipulate two observable matrices  $P_{T,H}$ ,  $P_H$ , and  $|A| \times |O|$  observable matrices  $P_{T,ao,H}$  Liu et al. (2016). Although Denis et al. Denis et al. (2014) showed that the concentration of the empirical  $P_{T,H}$  around its mean does not highly depend on its dimension, which gives a hope for alleviating the statistical problem when using large set of  $T$  and  $H$  as the accuracy of the learned model is directly connected to the concentration, manipulating these matrices is still too expensive to afford in large systems.

Thus, in practice, taking the computational constraints into account, it is needed to find only a finite set of columns of the Hankel matrix before the spectral methods can be applied. While different sets of columns usually lead to variant accuracy of the learned model, in this paper, we first introduce a concept of model entropy and show the high relevance between the entropy and the accuracy of the learned model, then for the columns selection problem, we call it basis selection in spectral learning of PSRs, an approach is proposed by using the model entropy as the guidance. We finally show the effectiveness of the proposed approach by executing experiments on *PocMan*, which is one benchmark domain in the literature.

We organize the remainder of this paper as follows. We briefly review the background and define notations in Section 2. We propose the approaches for basis selection in Section 3. We provide comparative results in Section 4. Finally we conclude the paper in Section 5.

## 2. Preliminaries

Predictive state representations (PSRs) represent state by using a vector of predictions of fully observable quantities (tests) conditioned on past events (histories), denoted  $b(\cdot)$ . For discrete systems with finite set of observations  $O = \{o^1, o^2, \dots, o^{|O|}\}$  and actions  $A = \{a^1, a^2, \dots, a^{|A|}\}$ , at time  $\tau$ , a *test* is a sequence of action-observation pairs that starts from time  $\tau + 1$ . Similarly, a *history* at  $\tau$  is a sequence of action-observation pairs that starts from the beginning of time and ends at time  $\tau$ , which is used to describe the full sequence of past events. The prediction of a length- $m$  test  $t$  at history  $h$  is defined as  $p(t|h) = p(ht)/p(h) = \prod_{i=1}^m Pr(o_i|ha_1o_1 \dots a_i)$  Singh and James (2004).

The underlying dynamical system can be described by a special bi-infinite matrix, called the Hankel matrix Balle et al. (2014), where the rows and columns correspond to all the possible tests  $T$  and histories  $H$  respectively, the entries of the matrix are defined as  $P_{t,h} =$

$p(ht)$  for any  $t \in T$  and  $h \in H$ , where  $ht$  is the concatenation of  $h$  and  $t$  Boots et al. (2011). The rank of the Hankel matrix is called the linear dimension of the system. When the rank is finite, we assume it is  $k$ , in PSRs, the state of the system at history  $h$  can be represented as a prediction vector of  $k$  tests conditioned at  $h$ . The  $k$  tests used as the state representation is called the minimal *core tests* that the predictions of these tests contain sufficient information to calculate the predictions for all tests, and is a *sufficient statistic* Singh and James (2004); Liu et al. (2015). For linear dynamical systems, the minimal core tests can be the set of tests that corresponds to the  $k$  linearly independent columns of the Hankel matrix Liu et al. (2015).

A PSR of rank  $k$  can be parameterized by a reference condition state vector  $b_* = b(\epsilon) \in \mathbb{R}^k$ , an update matrix  $B_{ao} \in \mathbb{R}^{k \times k}$  for each  $a \in A$  and  $o \in O$ , and a normalization vector  $b_\infty \in \mathbb{R}^k$ , where  $\epsilon$  is the empty history and  $b_\infty^T B_{ao} = 1^T$  Hsu et al. (2009); Boots et al. (2011). In the spectral approach, these parameters can be defined in terms of the matrices  $P_{\mathcal{H}}$ ,  $P_{\mathcal{T}, \mathcal{H}}$ ,  $P_{\mathcal{T}, ao, \mathcal{H}}$  and an additional matrix  $U \in \mathbb{R}^{T \times d}$  as shown in Eq. 1, where  $\mathcal{T}$  and  $\mathcal{H}$  are the set of all possible tests and histories respectively,  $U$  is the left singular vectors of the matrix  $P_{\mathcal{T}, \mathcal{H}}$ ,  $^T$  is the transpose and  $^\dagger$  is the pseudo-inverse of the matrix Boots et al. (2011).

$$\begin{aligned} b_* &= U^T P_{\mathcal{H}} 1_k, \\ b_\infty &= (P_{\mathcal{T}, \mathcal{H}}^T)^\dagger P_{\mathcal{H}}, \\ B_{ao} &= U^T P_{\mathcal{T}, ao, \mathcal{H}} (U^T P_{\mathcal{T}, \mathcal{H}})^\dagger. \end{aligned} \tag{1}$$

Using these parameters, after taking action  $a$  and seeing observation  $o$  at history  $h$ , the PSR state at next time step  $b(hao)$  is updated from  $b(h)$  as follows Boots et al. (2011):

$$b(hao) = \frac{B_{ao} b(h)}{b_\infty^T B_{ao} b(h)}. \tag{2}$$

Also, the probability of observing the sequence  $a_1 o_1 a_2 o_2 \cdots a_n o_n$  in the next  $n$  time steps can be predicted by Kulesza et al. (2015):

$$Pr[o_{1:t} | a_{1:t}] = b_\infty^T B_{a_n o_n} \cdots B_{a_2 o_2} B_{a_1 o_1} b_*. \tag{3}$$

Under the assumption that the columns and rows of the submatrix  $H_s$  of the Hankel matrix are sufficient, when more and more data are included, the law of large numbers guarantees that the estimates  $\hat{P}_{\mathcal{H}}$ ,  $\hat{P}_{\mathcal{T}, \mathcal{H}}$ , and  $\hat{P}_{\mathcal{T}, ao, \mathcal{H}}$  converge to the true matrices  $P_{\mathcal{H}}$ ,  $P_{\mathcal{T}, \mathcal{H}}$ , and  $P_{\mathcal{T}, ao, \mathcal{H}}$ , the estimates  $\hat{b}_*$ ,  $\hat{b}_\infty$ , and  $\hat{B}_{ao}$  converge to the true parameters  $b_*$ ,  $b_\infty$ , and  $B_{ao}$  for each  $a \in A$  and  $o \in O$ , that the learning is consistent Boots et al. (2011).

### 3. Basis Selection via Model Entropy

As the assumption that the columns and rows of the submatrix  $H_s$  are sufficient is really strong (usually means a very large set of columns and histories of the matrix), which almost fails in reality. At the same time, due to the computation and statistics constraints, we can only manipulate a limited finite set of tests/histories. As different sets of columns (tests) or rows (histories) usually cause different model accuracy, which requires us to select the finite set of tests and rows (histories) of the matrix for applying the spectral approach. In

practice, as all possible training data is used to generate the histories, how to select the tests, i.e., the basis selection, is the crucial problem. In this section, we first introduce the concept of model entropy, then show the high relevance between the entropy and the model accuracy, finally we propose a simple search method for selecting the bases.

**Model Entropy.** Given a set of tests  $X = \{x_1, x_2, \dots, x_i\}$ , if it includes the set of core tests and the number of possible  $p(X|\cdot)$  is finite, Proposition 1 holds Liu et al. (2016).

**Proposition 1** *The Markov decision process (MDP) model built by using  $p(X|\cdot)$  as state representation and the action-observation pair  $\langle ao \rangle$  as action is deterministic.*

In practice, it is difficult for  $X$  to include the set of core tests. At any time step, the prediction vector  $p(X|\cdot)$  may actually correspond to several PSR states. In such cases, the transition from  $p(X|h)$  to  $p(X|hao)$  usually becomes stochastic and the less information is included in  $X$ , the more stochastic the transition will usually be. Inspired by the concept of Shannon entropy that measures information uncertainty Shannon, to quantify the stochasticity, a concept of model entropy for each set of tests  $X$  is defined in Eq. 4 Liu et al. (2016).

$$E(X) = - \sum_{a \in \mathcal{A}_{PP}} \frac{1}{r(T^a)} \sum_{i=1}^{r(T^a)} \sum_{j=1}^{c(T^a)} T(s_i, a, s_j) \log T(s_i, a, s_j), \quad (4)$$

where  $T$  is the state-transition function of the MDP using  $p(X|\cdot)$  as state representation and  $\langle ao \rangle$  as action,  $\mathcal{A}_{PP} = A \times O$  the set of action-observation pairs in the original system,  $r(T^a)$  and  $c(T^a)$  the number of rows and columns of the state-transition matrix  $T^a$  respectively.

**Relevance between Model Entropy and Basis Selection.** When selecting the set of tests for applying spectral methods, the tests containing more information should be selected as the more information included in the set of tests, the higher accuracy the learned PSR model will usually be. According to Eq. 4, the less information is included in the set of tests  $X$ , the more stochastic the state transition will be, and the higher the model entropy is. So for the basis selection problem, the set of tests with lowest model entropy should be selected.

**Basis Selection via Model Entropy.** Using the model entropy as the guidance, we propose a simple local search algorithm for searching the set of tests used for spectral learning of PSRs, which is shown in Algorithm 1. Starting with a default  $T$  of the desired size, we iteratively sample a set of new tests and consider using it to replace the tests in  $T$ . If the replacement is a reduction in terms of the entropy value, then we keep it. After a fixed number of rounds, we stop and return the current  $T$ . In the algorithm, the entropy  $E(T)$  for each candidate set of tests  $T$  is calculated as follows (We name this procedure as *EntropyLearn*( $D, T$ )): We first generate an original randomly action-observation sequences  $D$  as the training data for calculating the entropy. Then the data is translated into the form of  $\langle \text{action-observation} \rangle \cdot p(T|\cdot)$  sequences. For example, a sequence  $d = \langle a_1 o_1 a_2 o_2 \dots a_k o_k \rangle$  is converted into  $d' = \langle a_1 o_1 \rangle p(T|a_1 o_1) \langle a_2 o_2 \rangle p(T|a_1 o_1 a_2 o_2) \dots \langle a_k o_k \rangle p(T|a_1 o_1 a_2 o_2 \dots a_k o_k)$ , where  $p(T|\cdot)$  can be estimated using the training data. Due to the sampling error, it is unlikely that any of these estimated  $p(\hat{T}|\cdot)$  will be exactly the same, even if the true underlying  $p(T|\cdot)$  are identical. Statistical tests or linearly independent techniques can be used to estimate the number of distinct underlying  $p(T|\cdot)$  and cluster the estimated  $p(\hat{T}|\cdot)$  corresponding to the same true prediction vector into one group(state) Talvitie and Singh

---

**Algorithm 1** Search for sets of  $k$  tests  $T$  that approximately minimize  $E(T)$

---

```

1: Input: dataset  $D$ , initial  $T$  of size  $k$ , number of size  $n < k$ , distributions  $p_T$  over
   candidate tests, number of rounds  $r$ , entropy value  $EV$ , an entropy threshold  $E_{th} > 0$ ,
   number of iteration  $iterNum$ .
2: Initialize:  $EV := EntropyLearn(D, T)$ .
3: for  $i = 1$  to  $r$  do
4:   for  $j = 1$  to  $iterNum$  do
5:     Sample a set of  $n$  tests  $T_s \notin T \sim p_T$ 
6:      $T' \leftarrow$  a set of  $n$  test in  $T$ 
7:      $EV_T := EntropyLearn(D, T \setminus T' \cup T_s)$ 
8:     if  $EV - EV_T > E_{th}$  then
9:        $EV := EV_T$ 
10:       $T := T \setminus T' \cup T_s$ 
11:      break
12:    end if
13:  end for
14: end for
15: Output:  $T$ 

```

---

(2011); Liu et al. (2014). Subsequently, we compute the state-transition functions in the transformed data and build the MDP model. Finally, the entropy  $E(T)$  of a set of tests  $T$  can be calculated according to Eq. 4.

## 4. Experimental Results

We evaluated the proposed technique on *PocMan* Silver and Veness (2010); Hamilton et al. (2014), a partially observable version of the video game *PacMan*. The environment is extremely large that has 4 actions, 1024 observations and an extremely large number of states (up to  $10^{56}$ ).

### 4.1. Experimental Setting

**Evaluated Methods.** We compared our approach (Entropy-based) with the bound-based approach Kulesza et al. (2015), which is the most related technique with our method that both the approaches address the same basis selection problem. The two approaches both first select a set of tests, then spectral method is applied by using these tests. Also, iterative approaches are used in both approaches for selecting the sets of tests. The main difference between our approach and the work of Kulesza et al. (2015) is the guidance used for selecting the set of tests. While our approach uses the model entropy as the guidance, the key idea in Kulesza et al. (2015) is that as in the limiting-case, the singular values of the learned transformed predictive state representations (TPSR) parameters are bounded, then under an assumption that the smaller the singular values of the learned parameters are, the more accuracy the learned PSR model will be, they just select the tests that can cause smaller singular values of the learned PSR parameters. However, no any formal guarantees are provided for such an approach.

A uniformly random generated sequence with length 200,000 was used as the training sequence for both approaches. And to calculate the entropy of a candidate set of tests  $T_c$ , a randomly generated 5,000-length action-observation sequence  $d$  was used. For each test  $t \in T_c$ ,  $p(t|\cdot)$  was estimated by executing  $act(t)$  100 times, where  $act(t)$  is the action sequence in test  $t$ . The entropy thresholds for our approaches are 0.06, 0.04, and 0.02 for different number of tests as the entropy value usually decreases with the increase of the number of tests (For number of tests 100 and 150, the threshold used is 0.06; for number of tests 200 and 250, the threshold is 0.04; and for number of tests 300, the threshold is 0.02). The number of rounds for both approaches is 10. In each round, we iteratively sample a set of new tests of size 20 and use it to replace 20 tests in  $T$ , the iterative number in each round is also 10. After a fixed number of rounds, we stop and return the current  $T$ .

**Performance Measurements.** We evaluated the learned models in terms of prediction accuracy, which is measured by the difference between the true predictions and the predictions given by the learned model over a test data sequence (for *PocMan*, we cannot obtain the true predictions, Monte-Carlo rollout predictions were used [Hamilton et al. \(2014\)](#)).

Two error functions were used in the measurement. One is the average one-step prediction error per time step on the test sequence as shown in Eq. 5:

$$\frac{1}{L} \sum_{t=1}^L |p(o_{t+1}|h_t, a_{t+1}) - \hat{p}(o_{t+1}|h_t, a_{t+1})|. \quad (5)$$

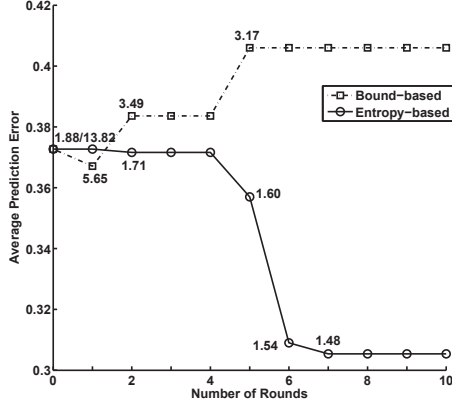
$p(\cdot)$  is the probability calculated from the true POMDP model or the Monte-Carlo rollout prediction and  $\hat{p}(\cdot)$  the probability obtained from the learned model.  $|\cdot|$  refers to the absolute value.  $L$  is the length of the test sequence used in the experiments. For the average four-step prediction error, the same equation with predicting four steps ahead  $\hat{p}(o_{t+1}o_{t+2}o_{t+3}o_{t+4}|h_t a_{t+1} a_{t+2} a_{t+3} a_{t+4})$  was used.

## 4.2. Performance Evaluation

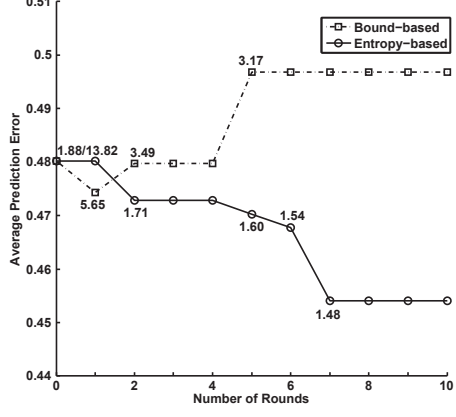
In this section, for each algorithm, we report the performance results as the mean error over 10 trials. For each trial, a uniform randomly generated test sequence with length  $L = 20,000$  was used for testing the accuracy of the learned model.

Two kinds of experiments were conducted. The first experiment is to evaluate the prediction performance by fixing the number of tests as 100, and ran Algorithm 1 10 rounds. For each round, we reported the average one and four prediction errors for both approaches, the results are shown in Fig. 1. The number at each point in Fig. 1 is the average entropy (for our approach) or the average largest singular values (for the bound-based approach) in the corresponding round, which also shows that for the entropy-based approach, a higher entropy value results in a lower prediction accuracy while there are no relevance between the singular values and the prediction accuracy. The second experiment is to evaluate the prediction performance by varying the number of tests. For each number of tests, both approaches were ran 10 rounds, and the final results for both one and four-step errors after 10 rounds were reported in Fig. 2. For this experiment, we also reported the initial results (Initial) without the replacement of the tests.

As can be seen from Figs. 1 and 2, both for the one-step and four-step predictions, and both for the two kinds of experiments, in all the cases, our algorithm performs very

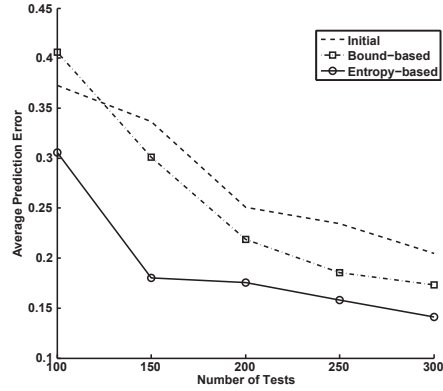


(a)

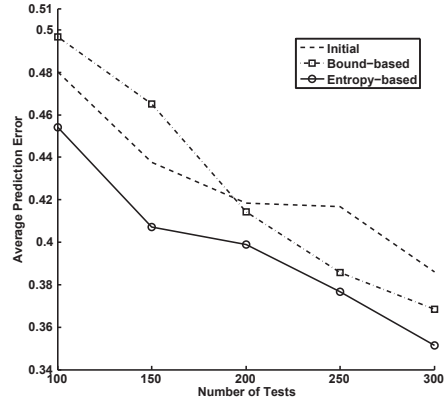


(b)

Figure 1: (a) One-step; (b) four-step prediction error of 10 rounds for fix number of tests



(a)



(b)

Figure 2: (a) One-step; (b) four-step prediction error of different number of tests



well and outperforms the bound-based approaches. Considering the error reported is only the prediction error for one time step, the improvement on prediction accuracy on a long sequence is remarkable. Meanwhile, for the first kind of experiment, as the number of rounds increases, our algorithm reduces its prediction error while the bound-based approaches with increasing rounds do not improve their performances and are very unstable. What also can be seen from the experimental results is that with the increase of the number of tests, for all the approaches, the prediction accuracy of the obtained PSR model also increases, which demonstrates the importance of including more tests in the learning of the models.

## 5. Conclusion

How to choose the bases is a very important problem in spectral learning of PSRs. However, until now, there are very little work that can address this issue successfully. In this paper, by introducing the model entropy for measuring the model accuracy and showing the close relevance between the entropy and model accuracy, we propose an entropy-based basis selection strategy for spectral learning of PSRs. Several experiments were conducted on the PocMan environment, and the results show that compared to the state-of-the-art bound-based approach, our technique is more stable and achieved much better performance.

## Acknowledgments

This work was supported by the National Natural Science Foundation of China (No. 61375077).

## References

- Borja Balle, Xavier Carreras, Franco M. Luque, and Ariadna Quattoni. Spectral learning of weighted automata. *Mach. Learn.*, 96(1-2):33–63, July 2014. ISSN 0885-6125. doi: 10.1007/s10994-013-5416-x. URL <http://dx.doi.org/10.1007/s10994-013-5416-x>.
- Byron Boots, Sajid Siddiqi, and Geoffrey Gordon. Closing the learning planning loop with predictive state representations. *International Journal of Robotic Research*, 30:954–956, 2011.
- François Denis, Mattias Gybels, and Amaury Habrard. Dimension-free concentration bounds on hankel matrices for spectral learning. In *Proceedings of the 31th International Conference on Machine Learning, ICML*, pages 449–457, 2014. URL <http://jmlr.org/proceedings/papers/v32/denis14.html>.
- William Hamilton, Mahdi Milani Fard, and Joelle Pineau. Efficient learning and planning with compressed predictive states. *Journal of Machine Learning Research*, 15:3395–3439, 2014. URL <http://jmlr.org/papers/v15/hamilton14a.html>.
- William L. Hamilton, Mahdi M. Fard, and Joelle Pineau. Modelling sparse dynamical systems with compressed predictive state representations. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, pages 178–186, 2013.
- Daniel Hsu, Sham M. Kakade, and Tong Zhang. A spectral algorithm for learning hidden markov models. In *The 22nd Annual Conference on Learning Theory (COLT 2009)*, 2009.



- Alex Kulesza, Nan Jiang, and Satinder Singh. Spectral learning of predictive state representations with insufficient statistics. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, 2015.
- Michael L. Littman, Richard S. Sutton, and Satinder Singh. Predictive representations of state. In *Advances In Neural Information Processing Systems (NIPS)*, pages 1555–1561, 2001.
- Yunlong Liu, Zijiang Yang, and Guoli Ji. Solving partially observable problems with inaccurate psr models. *Information Sciences*, 283:142–152, 2014.
- Yunlong Liu, Yun Tang, and Yifeng Zeng. Predictive state representations with state space partitioning. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1259–1266, 2015.
- Yunlong Liu, Hexing Zhu, Yifeng Zeng, and Zongxiong Dai. Learning predictive state representations via monte-carlo tree search. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.
- Matthew Rosencrantz, Geoff Gordon, and Sebastian Thrun. Learning low dimensional predictive representations. In *Proceedings of the Twenty-first International Conference on Machine Learning (ICML)*, pages 695–702, 2004.
- Claude Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, 623–656, 1948.
- David Silver and Joel Veness. Monte-carlo planning in large pomdps. In *In Advances in Neural Information Processing Systems 23 (NIPS)*, pages 2164–2172, 2010.
- Satinder Singh and Michael R. James. Predictive state representations: A new theory for modeling dynamical systems. In *Proceedings of the Twentieth Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 512–519, 2004.
- Erik Talvitie and Satinder P. Singh. Learning to make predictions in partially observable environments without a generative model. *J. Artif. Intell. Res. (JAIR)*, 42:353–392, 2011.